

A Gaussian sliding windows regression model for hydrological inference

Stefan Schrunner^{1,2} , Parham Pishrobat³ , Joseph Janssen⁴ ,
Anna Jenul¹ , Jiguo Cao⁵ , Ali A. Ameli⁴  and William J. Welch³ 

¹Department of Data Science, Norwegian University of Life Sciences, Universitetstunet 3, 1433 Ås, Norway

²Carinthia University of Applied Sciences, Europastraße 4, 9524 Villach, Austria

³Department of Statistics, University of British Columbia, 2207 Main Mall, Vancouver, BC, Canada, V6T 1Z4

⁴Department of Earth, Ocean and Atmospheric Sciences, University of British Columbia, 2207 Main Mall, Vancouver, BC, Canada, V6T 1Z4

⁵Department of Statistics and Actuarial Science, Simon Fraser University, 8888 University Drive, Burnaby, BC, Canada, V5A 1S6

Address for correspondence: Stefan Schrunner, Carinthia University of Applied Sciences, Europastraße 4, 9524 Villach, Austria. Email: s.schrunner@fh-kaernten.at

Abstract

Statistical models are an essential tool to model, forecast, and understand the hydrological processes in watersheds. In particular, the understanding of time lags associated with the delay between rainfall occurrence and subsequent changes in streamflow is of high practical importance. Since water can take a variety of flow paths to generate streamflow, a series of distinct runoff pulses may combine to create the observed streamflow time series. Current state-of-the-art models are not able to sufficiently confront the problem complexity with interpretable parametrization, thus preventing novel insights about the dynamics of distinct flow paths from being formed. The proposed Gaussian Sliding Windows Regression Model targets this problem by combining the concept of multiple windows sliding along the time axis with multiple linear regression. The window kernels, which indicate the weights applied to different time lags, are implemented via Gaussian-shaped kernels. As a result, straightforward process inference can be achieved since each window can represent one flow path. Experiments on simulated and real-world scenarios underline that the proposed model achieves accurate parameter estimates and competitive predictive performance, while fostering explainable and interpretable hydrological modelling.

Keywords: flow paths, Gaussian kernel, interpretability, sliding windows regression, streamflow prediction

1 Introduction

The hydrological processes that produce streamflow play key roles in determining the environmental effects of climate and land use changes. In particular, changes in climate or land use can trigger a complex series of nonlinear and interactive processes which can eventually impact the way in which watersheds partition, store, and release water, leading to potential changes in flood, landslide or drought risks (Dunn et al., 2010; Harman et al., 2011; Sawicz et al., 2014). Furthermore, water reservoirs and consistent streamflow are essential for the regulation of the water supply in both urban and rural areas (Janssen et al., 2021; Tang et al., 2009), as well as for sustainable energy production and continually healthy ecological habitats (Ahmad & Hossain, 2020; Zalewski, 2000). A deep understanding of how precipitated water becomes streamflow is of great importance and requires a sophisticated and interpretable statistical framework.

Hence we investigate the problem of inferring streamflow partitioning into different flow paths such as overland flow, subsurface flow, and baseflow (Cornette et al., 2022; Kannan et al., 2007;

Received: October 4, 2023. Revised: January 20, 2025. Accepted: January 22, 2025

© The Royal Statistical Society 2025.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

McMillan, 2020; Nejadhashemi et al., 2009; Wang & Gupta, 2024). While there exist many simple methods which can partition streamflow into fast surface flow versus slow baseflow (Nejadhashemi et al., 2009; Wang & Gupta, 2024), the processes and reasoning behind this partitioning are usually lost. To prevent this, one could consider how input rainfall is partitioned into different flow paths. Hence, we consider a time series $y_T = (y_t)_{t \in T}$ of streamflow at a particular gauge and an input time series $x_T = (x_t)_{t \in T}$ of rainfall data on a common time domain T . Given such data, a common approach is a lagged regression model following the concepts of auto-regressive moving average with exogenous variables (ARIMAX, ARMAX) (Hyndman & Athanasopoulos, 2021). In contrast to these models, our problem does not involve an auto-regressive part, i.e. we do not use lagged observations of the target y_t as predictors; instead, we seek to fully characterize the target time series via the lagged input time series. Further, we desire more interpretability and inferential power than is offered by ARIMAX models. For example, the maximum time interval for which the predictor has an impact on the response and the separation of flow pathways is of interest for watershed modelling.

A more closely related model in time series analysis is the finite distributed lag model (DLM) (Baltagi, 2022). Despite its roots in econometrics, the DLM has been extensively applied in environmental sciences (Chen et al., 2018; Peng et al., 2009; A. M. Rushworth et al., 2013; Warren et al., 2020). While ARIMAX mainly relies on the concept of ARIMA, i.e. it makes use of the lagged target variable y_t as input, and only adds the present time point of the exogenous variable x_t , DLM predicts the target variable purely based on a distinct lagged input time series—the same setup that is pursued in this work. In this sense, the presented work may be seen as a specific type of a DLM and thereby an extension of the work by A. M. Rushworth et al. (2013), however, with an improved parameterization, which allows for improved model interpretability. Basic DLM variants do not allow us to impose parametric assumptions on the shape of the regression parameter vector beyond simple singular shapes like geometrically decreasing patterns or polynomials (Almon, 1965; Eisner, 1960; Griliches, 1967). Such DLM extensions permit a large variety of lag curve shapes, but do not make the separation between flow paths explicit. To the best of our knowledge, none of the existing DLM variants are capable of explicitly separating distinct flow pulses, which is most crucial for hydrologic inference.

There are several other possible approaches. Another economic model was adapted for hydrological inference by Giani et al. (2021), where the authors deploy the concept of cross-correlations in combination with a moving average approach. Although their method has favourable properties to model the response time between rainfall events and flow pulses, it again does not allow a direct separation of the flow paths. Furthermore, correlation-based analyses do not allow straightforward predictions on new test datasets. Similarly, we do not pursue continuous-time models in hydrology based on partial differential equations (e.g. Young, 2006).

In contrast to further related works in the field of statistical modelling in hydrology, the methodology presented here focuses exclusively on the temporal relation between rainfall and streamflow in a given watershed. Spatial relationships between watersheds, as investigated by Roksvåg et al. (2021), are beyond the scope of this article; nevertheless, the presented concept offers a foundation, which can be extended to include spatial information for gauged and ungauged catchments in future work and may complement existing large-scale analyses, such as that of Hare et al. (2021).

In this article, we present a parameterized, interpretable variant of lagged regression models for hydrological modelling: the Gaussian sliding windows regression (SWR) model. It builds on the assumption that distinct flow paths (surface flow, subsurface flow, etc.) provoke separable pulses of streamflow after rainfall events. These pulses are represented by multiple temporal windows, which weight the input x_t at time points within the window via a Gaussian kernel. The output of each window is mapped to the target variable y_t using a linear regression model. Compared to earlier SWR models (Davtyan et al., 2020; Janssen et al., 2021; Khan et al., 2019), the proposed method has major structural differences, which constitute the main novelties of our work:

1. We consider multiple, potentially overlapping windows representing distinct flow paths, and
2. we use a parameterized kernel as model weights to allow straightforward interpretations for hydrological inference of flow path importance.

Our experimental evaluation involves both simulated and real-world data. The simulation study serves as a proof of concept for the model and validates the parameter estimation procedure. The application to two real-world datasets demonstrates that the model achieves competitive predictive performance and has favourable properties for interpreting underlying hydrological processes.

2 SWR model

We assume that a univariate input time series $\mathbf{x}_T$ and a univariate target time series $\mathbf{y}_T$ are given on a common, discrete time domain $T = [t_{\max}]$, where $[n]$ denotes the index set $\{0, \dots, n\}$. In the motivating application $\mathbf{x}_T$ and $\mathbf{y}_T$ represent daily rainfall and streamflow, respectively. The modelling goal is to represent the target value y_t at time index $t \in T$ based on time-lagged observations of $\mathbf{x}_{[t]}$, i.e. up to time t .

Our proposed model pursues the idea that the effect of precipitation on streamflow follows a mixture of k kernels associated with temporal windows, each one characterizing a particular flow path (groundwater flow, subsurface flow, or overland flow). Ultimately, we aim to represent the distribution of response times from flow path $i \in \{1, \dots, k\}$ by weight vector $\boldsymbol{\kappa}^{(i)} \in \mathbb{R}^{t+1}$. Note that the length of $\boldsymbol{\kappa}^{(i)}$ is formally time-dependent based on its dimension. Since the subvector of nonzero weights remains the same over time, zero padding is used at the end to comply with the dimensionality of $\mathbf{x}_{[t]}$. Thus, the time index is left out in a slight abuse of notation. Computing a weighted sum of the lagged time series $\mathbf{x}_{[t]}$ with one such weight vector $\boldsymbol{\kappa}^{(i)}$ models the contribution to streamflow y_t associated with flow path i . Summing the contributions gives the overall flow. To facilitate explanations, we present the model for the special case of only $k = 1$ window first, and extend it to multi-window scenarios afterwards.

2.1 Single-window model

If only one flow path exists, we denote the associated weight vector by $\boldsymbol{\kappa} \in \mathbb{R}^{t+1}$. By default, we require that all entries in $\boldsymbol{\kappa}$ are nonnegative, and that $\boldsymbol{\kappa}$ is normalized, i.e. $\|\boldsymbol{\kappa}\|_1 = 1$. Then, in the special case of one window, our SWR model describes the gauged runoff y_t by the weighted sum or the associated convolution

$$\begin{aligned} y_t &= \beta \cdot \left(\sum_{s=0}^t x_{t-s} \kappa_s \right) + \varepsilon_t \\ &= \beta \cdot (\mathbf{x}_{[t]} * \boldsymbol{\kappa}) + \varepsilon_t, \end{aligned} \quad (1)$$

with some multiplicative constant β and an error term ε_t assumed as Gaussian white noise. The symbol $*$ denotes the discrete convolution operator defined as

$$\mathbf{v} * \mathbf{z} = \sum_{i=0}^n v_{n-i} z_i$$

for two vectors $\mathbf{v}, \mathbf{z} \in \mathbb{R}^{n+1}$. Note that weight κ_s is applied to lag $s \in [t]$, i.e. κ_s denotes the weight applied to observation x_{t-s} rather than x_s . The model parameter β acts as a regression coefficient to adjust the scaling of the normalized weight vector $\boldsymbol{\kappa}$. Since additional precipitation cannot lead to decreased streamflow, we further require $\beta > 0$. In other applications, this restriction may be modified.

Unlike typical linear regression models, the SWR model is defined without an intercept $\beta^{(0)}$. This is due to the assumption that runoff y_t is 0 after a long period without precipitation, $x_s = 0$ for all $s < t$. However, in other applications, the proposed SWR model may be used with an intercept.

Further, note that the model definition allows negative responses y_t if all entries of β are very small while the error ε_t takes a larger negative value. However, in this rare case, point forecasts remain nonnegative due to nonnegative observations $\mathbf{x}_{[t]}$.

The assumption of independent errors ε_t in model (1) (or its multi-kernel generalization) is unrealistic for most real-world datasets, where autocorrelated disturbances are common. Later, we describe a simple data transformation to address that complexity, however, it is sufficient to proceed with the simpler model (1) for now.

The proposed SWR model (1) is related to a DLM, which is a common approach also used in related work (A. M. Rushworth et al., 2013). However, the paradigm where each flow path is represented by one kernel κ allows us to make further assumptions to parametrize the lag vector: We assume that each flow path is associated with a distinct unimodal distribution of the time lag (time delay) between a rainfall event and its effect on streamflow. Hence, the weight vector κ represents the kernel of a probability distribution. Further, we denote the time lag with the highest amount of runoff water, represented by the position of the peak of κ (mode), as a distribution centre parameter δ , where $0 < \delta < t$.

More specifically, we parametrize κ by a Gaussian-shaped kernel centred around a lag of δ , assuming that the time delay between the rainfall event and the resulting change in the streamflow is approximately Gaussian distributed. Since the time series is only measured at discrete time points, the Gaussian kernel is discretized by applying a step function, such that

$$\kappa_s = \frac{1}{C} \int_{s-\frac{1}{2}}^{s+\frac{1}{2}} \phi(\tau; \delta, \sigma^2) d\tau,$$

where $\phi(\cdot, \delta, \sigma^2)$ denotes the probability density function of a Gaussian distribution with mean δ and variance $\sigma^2 > 0$, the kernel indices (lags) $s \in \{0, 1, \dots, t\}$ are relative to a current time point, and C denotes the normalization constant

$$C = \sum_{s=0}^t \int_{s-\frac{1}{2}}^{s+\frac{1}{2}} \phi(\tau; \delta, \sigma^2) d\tau. \quad (2)$$

Note that the first kernel element, κ_0 , corresponding to lag $s = 0$ is supported on $(-\frac{1}{2}, \frac{1}{2})$, i.e. including half a time unit before time t , because of the discretization.

The notion of a *temporal window* describes an interval of time lags $\delta \pm r$ with some window size $2r$ centred around δ , which covers a large proportion of the probability mass. Due to the shape of the Gaussian probability density function and the decay when moving away from δ , we define the window such that it covers approximately a range of ± 3 standard deviations, i.e. $r = 3\sigma$. For now, assume that $\delta - r > 0$ and $\delta + r < t_{\max}$ (the special cases where the boundaries of $[t]$ are exceeded will be treated at a later point). Given such a window, the modelled response y_t exclusively depends on $\mathbf{x}_{[t]}$ falling in $([t - \delta - r], [t - \delta + r])$ rather than the full index set $[t]$. Here, $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ denote floor and ceiling operators to consider the weights of the two discrete time points just outside the kernel. Nevertheless, to keep the notation manageable, we formally retain the notation of κ covering the whole interval $[t]$ in the following.

Examples for $r \in \{1, 3, 6\}$ and $\sigma = r/3$ are depicted in Figure 1. Further, the graphical idea behind the SWR model for a univariate input $\mathbf{x}$ can be visualized in Figure 2: the target time series y_t at time t is predicted based on information originating mainly from the input time series in window $\delta \pm r$.

In the special case that σ is very small and r approaches 0, the window width collapses to 0, and the full probability mass of the kernel accumulates at the window centre δ . However, since δ may be located between two observations of the input time series, σ is restricted to be at least $\frac{1}{6}$ to ensure at least one observation falls within the effective domain of the probability mass, i.e. is assigned to the closest measured time point in such cases, by definition.

Nonzero weights can only be assigned to lags in $[t]$, and it follows that the Gaussian SWR model cannot be evaluated for arbitrary parameter combinations of δ and σ . The restriction on the upper (lag) limit of the window, i.e. $\delta + r \leq t$, is less prohibitive in practice for longer time series and merely leads to y_t not being predictable for $t < \delta + r$, i.e. for small t at the beginning of the time series. As a consequence, this restriction can be easily resolved by withholding data for calibration during model training. Hydrological applications typically have time lags δ in the range of a few days, weeks, or months, hence the impact of this restriction is negligible for a sufficiently long time series spanning over multiple years.

On the other hand, peaks appearing at short time lags are highly relevant in practice, leading to small values of δ . Under extremely wet conditions in low permeability catchments, it can be




Figure 1. Examples of Gaussian-shaped window kernels centred around $t - \delta$ with $\delta = 7.3$. For different values of σ a window extends approximately $\pm r$ around $t - \delta$, where $r = 3\sigma$. Each bullet denotes a discrete time point of lag s with weight given by the area under the kernel between $t - s - 1/2$ and $t - s + 1/2$. A lagged time point at $t - s$ outside $t - \delta \pm r$, such as those at $t - 6$ (top), $t - 4$ (middle), and $t - 1$ (bottom), may still have a small weight as $t - s \pm 1/2$ falls within the kernel. Other time points immediately outside their respective kernels, such as those at $t - 9$ (top), $t - 11$ (middle), and $t - 14$ (bottom), will be considered but have essentially zero weight as $t - s \pm 1/2$ does not reach nontrivial values of the Gaussian kernel.




Figure 2. Illustration of the Sliding Windows Regression model with $k = 1$ window, which predicts y_{70} based on a subset of $\mathbf{x}_{[70]}$ values in a window. With $\delta = 20$ and $r = 10$, the window $[40, 60]$ centred around $s - \delta = 50$ covers $2r + 1 = 21$ time points.




Figure 3. Example of a truncated Gaussian-shaped window kernel, with window parameters $\delta = 2.7$ and $r = 3\sigma = 6$. The window spanning $t - \delta \pm r$ shown by the dashed rectangle is truncated to the shaded region: rainfall after time t cannot contribute to y_t . Hence, the weights for x_{t+1} , x_{t+2} , and x_{t+3} in the time interval marked τ are all zero even though the interval is within the rectangle. The Gaussian curve is renormalized to 1. The kernel weight applied to x_t is the area under the curve between $t - \frac{1}{2}$ and $t + \frac{1}{2}$ due to discretization.

assumed that water from rainfall will almost immediately lead to changes in streamflow. If the lower limit of the lag window, $\delta - r$, falls below 0, we change the shape of the kernel to a truncated Gaussian kernel, as shown in Figure 3. Note that the kernel is, by definition in (2) always normalized.

Truncation has effects on the interpretation of the model parameters: while δ still indicates the time lag with peak kernel weight (mode of the probability distribution), it does not indicate the central lag of the window on the time axis anymore. Further, the right-truncated Gaussian distribution is no longer symmetric with respect to the temporal window, and hence, the distribution mean shifts to the left.

2.2 Multi-window model

Assuming multiple different (potentially overlapping) windows allows the model to account for distinct effects on the current target y_t . For this purpose, we extend our method from a single-window model to a mixture of k kernels, which is related to the density resulting from a mixture of probability distributions (Everitt et al., 1981). For a given window $i \in \{1, \dots, k\}$, we use the same notation as in the single-window case for the window and kernel parameters with an added index i denoting the window number: $\delta^{(i)}$, $r^{(i)}$, and $\sigma^{(i)}$. The Gaussian SWR model with multiple windows is given as a direct generalization of (1) by

$$y_t = \sum_{i=1}^k \beta^{(i)} \cdot (\mathbf{x}_{[t]} * \boldsymbol{\kappa}^{(i)}) + \varepsilon_t \quad (3)$$

$$= \mathbf{x}_{[t]} * \left(\sum_{i=1}^k \beta^{(i)} \boldsymbol{\kappa}^{(i)} \right) + \varepsilon_t, \quad (4)$$

where the regression parameters $\beta^{(1)}, \dots, \beta^{(k)} \geq 0$ act as window weights. In summary, the Gaussian SWR model with k windows contains the following parameters:

- k regression parameters $\boldsymbol{\beta} = (\beta^{(1)}, \dots, \beta^{(k)}), \beta^{(i)} \in \mathbb{R}^+$,
- k lag parameters $\boldsymbol{\delta} = (\delta^{(1)}, \dots, \delta^{(k)}), \delta^{(i)} \in \mathbb{R}^+$, and
- k kernel standard deviation (size) parameters $\boldsymbol{\sigma} = (\sigma^{(1)}, \dots, \sigma^{(k)}), \sigma^{(i)} \in \mathbb{R}^+$.

The concept of a Gaussian SWR model with multiple windows is illustrated in Figure 4: a target value y_t is predicted by multiple windows accumulating information in $\mathbf{x}_{[t]}$ over different time intervals. As depicted in Figure 5, the combined kernel represents a linear combination of the single-window kernels $\boldsymbol{\kappa}^{(i)}$ for $i = 1, \dots, k$. Under the assumption that each window has a distinct centre point $\delta^{(i)}$ (otherwise, identifiability of the parameters might not be given), without loss of generality, we can order the location parameters such that $\delta^{(1)} < \delta^{(2)} < \dots < \delta^{(k)}$.




Figure 4. Illustration of the Sliding Windows Regression model with $k = 2$ windows, which predicts y_{70} based on $\mathbf{x}_{[70]}$. The window location parameters are $\delta^{(1)} = 20$ and $\delta^{(2)} = 40$. On the time axis, the window $[40, 60]$ is centred around $s - \delta^{(1)} = 50$ and covers $r^{(1)} = 10$ time points on each side of $\delta^{(1)} = 50$, while window $[10, 50]$ is centred around $s - \delta^{(2)} = 30$ and covers $r^{(2)} = 20$ time points on each side.




Figure 5. Example mixture of window kernels (weights): kernel $\kappa^{(1)}$ represents a truncated Gaussian, while kernel $\kappa^{(2)}$ has a Gaussian shape. The combined window $\beta^{(1)}\kappa^{(1)} + \beta^{(2)}\kappa^{(2)}$ is a mixture of Gaussians for the SWR model.

2.2.1 Parameter interpretation

The number of windows k models the number of underlying, separable flow paths. Since single-window kernels κ are normalized, regression parameters (*window coefficients*) $\beta^{(1)}, \dots, \beta^{(k)}$ act as weights among the kernels. For practical evaluations, they may be converted into proportions representing the relative impacts of the underlying flow paths.

For a window that is essentially untruncated, the location parameter $\delta^{(i)}$ and the width $\sigma^{(i)}$ of kernel i are loosely related to the mean and the standard deviation, respectively, of the time lag between precipitation and a change in streamflow. In any case $\delta^{(i)}$ is always the window mode and is thus a direct estimate of the typical time delay between rainfall events and the resulting impact on the gauged runoff for a flow path.

2.3 Information criteria

Based on the definition and the model assumptions of the proposed Gaussian SWR model, the model errors are given by

$$\varepsilon_t = y_t - \mathbf{x}_{[t]} * \left(\sum_{i=1}^k \beta^{(i)} \kappa^{(i)} \right),$$

and follow a Gaussian white noise process. Hence, the log-likelihood $\ln L(\boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\sigma}; \mathbf{x}_T, \mathbf{y}_T)$ is tractable and matches that of ordinary multiple linear regression models.

To choose the number of windows, we use the Bayesian Information Criterion (BIC, Schwarz, 1978). The number of model parameters in a k -window model aggregates one regression parameter, one window location parameter, and one window width parameter for each of the k windows. Hence,

$$\text{BIC} = -2 \ln L(\boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\sigma}; \mathbf{x}_T, \mathbf{y}_T) + \ln(|T|) \cdot 3k, \quad (5)$$

where $|T|$ is the number of observed time points (number of samples).

2.4 Model training and implementation

In a training dataset, let $\mathbf{y}_T$ denote the observed values of the target time series, and let $\hat{\mathbf{y}}_T(\boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\sigma})$ denote the vector of predictions as a function of the model parameters; the likelihood follows from $\mathbf{y}_T$ and $\hat{\mathbf{y}}_T$. Since established approaches to fitting lagged regression models, such as Hannan (1967) are not directly applicable due to our specific parametrization, Algorithm 1 outlines an iterative procedure to train the model by minimizing the negative log-likelihood (or another loss function). In each iteration $k = 1, \dots, k_{\max}$, one new window is added to the model and the parameters for all k windows are optimized. An information criterion (BIC for all simulations and real-data analyses) chooses the best number of windows among $1, \dots, k_{\max}$.

We use the GENOUD algorithm (Mebane & Sekhon, 2011) in the R package `rgenoud` (Mebane et al., 2023) to train a model with a fixed number of windows. GENOUD optimizes a population of candidate solutions over a number of generations via operations analogous to

Algorithm 1 Train a Gaussian SWR model, choosing the number of windows

Input:

$\mathbf{x}_T, \mathbf{y}_T$ (time series supported on a common set of times T)

$k_{\max}$ (maximum number of windows)

n (population size for genetic optimization)

$\mathbf{D}^{(0)}$ (a vector of starting values for the $\delta^{(i)}$)

$\mathbf{S}^{(0)}$ (a vector of starting values for the $\sigma^{(i)}$)

Set window counter $k \leftarrow 1$

Compute m_x and m_y , the mean of $\mathbf{x}_T$ and $\mathbf{y}_T$, respectively

for $k = 1$ to $k_{\max}$ **do**

if $k = 1$ **do**

 Generate an $n \times 1$ column vector of starting values for each parameter in window 1

$\mathbf{B}^{(1)} \leftarrow (m_y/m_x, \dots, m_y/m_x)^\top$ (for $\beta^{(1)}$)

$\mathbf{D}^{(1)} \leftarrow (\mathbf{D}^{(0)}, \mathbf{D}^{(0)}, \dots)^\top$ (for $\delta^{(1)}$)

$\mathbf{S}^{(1)} \leftarrow (\mathbf{S}^{(0)}, \mathbf{S}^{(0)}, \dots)^\top$ (for $\sigma^{(1)}$)

else

 Generate an $n \times 1$ vector of starting values for $\beta^{(1)}, \dots, \beta^{(k)}$

for $j = 1$ to k **do**

$\mathbf{B}^{(j)} \leftarrow (m_y/m_x/k, \dots, m_y/m_x/k)$ (for $\beta^{(j)}$)

end for

 Generate an $n \times 1$ vector of starting values for $\delta^{(k)}$ (note $\mathbf{D}^{(1)}, \dots, \mathbf{D}^{(k-1)}$ are unchanged)

$\mathbf{D}^{(k)} \leftarrow \text{permute}(\mathbf{D}^{(1)})$

 Generate an $n \times 1$ vector of starting values for $\sigma^{(k)}$ (note $\mathbf{S}^{(1)}, \dots, \mathbf{S}^{(k-1)}$ are unchanged)

$\mathbf{S}^{(k)} \leftarrow \text{permute}(\mathbf{S}^{(1)})$

end if

 Generate an $n \times (3k)$ population of starting values with columns for $\beta^{(1)}, \dots, \beta^{(k)}, \delta^{(1)}, \dots, \delta^{(k)}, \sigma^{(1)}, \dots, \sigma^{(k)}$

$$P = (\mathbf{B}^{(1)} \mid \dots \mid \mathbf{B}^{(k)} \mid \mathbf{D}^{(1)} \mid \dots \mid \mathbf{D}^{(k)} \mid \mathbf{S}^{(1)} \mid \dots \mid \mathbf{S}^{(k)})$$

 Minimize the negative log-likelihood for k windows via the GENOUD algorithm with starting values in P

$$(\hat{\beta}^{(k)}, \hat{\delta}^{(k)}, \hat{\sigma}^{(k)}) \leftarrow \arg \min_{\beta, \delta, \sigma} -\ln L(\beta, \delta, \sigma; \mathbf{x}_T, \mathbf{y}_T).$$

 Compute BIC_k for k windows from (5)

end for

 return $\hat{k}$ and $\hat{\beta}^{(\hat{k})}, \hat{\delta}^{(\hat{k})}, \hat{\sigma}^{(\hat{k})}$, such that information criterion BIC_k is optimized

$$\hat{k} = \arg \min_{k \in \{1, \dots, k_{\max}\}} \text{BIC}_k$$

genetic evolution. It also fine-tunes the best candidate found at each generation with a quasi-Newton method. Trial and error comparisons with a more standard-constrained optimizer (BOBYQA, see [Powell, 2009](#)) demonstrated a clear advantage in the optimal log-likelihood found, which is important if BIC is to reliably choose the number of windows. As a stopping criterion, the absolute tolerance is set to 10^{-3} , which is a trivial change in the log-likelihood.

In the GENOUD algorithm, we constrain the parameters such that all $\delta^{(i)} > 0$ and all $\sigma^{(i)} > 1/6$, which allows for truncated as well as nontruncated kernels. In the boundary case of $\delta = 0$, a

maximum level of truncation is reached for the Gaussian kernel. The constraint $\sigma^{(i)} > 1/6$, avoids collapse of the window. Moreover, the minimum window width is hence 1, i.e. $\pm 3 \times 1/6$, which is commensurate with the limitation of discrete data at unit time steps.

GENOUD requires a population of n starting values for β , δ , and σ . Experimental observations showed that this initial population is important. It needs to provide coverage of reasonable values in the parameter space but not be so overly wide as to make the optimization inefficient. Note that the initial ranges can be exceeded later during the GENOUD iterations. The population formed in Algorithm 1 has n members in the rows of a matrix ($n = 100$ for all results reported) and $3k$ columns for initial values of $\beta^{(1)}, \dots, \beta^{(k)}, \delta^{(1)}, \dots, \delta^{(k)}, \sigma^{(1)}, \dots, \sigma^{(k)}$, respectively. For $k = 1$, $\beta^{(1)}$ is always set to m_y/m_x , where m_x and m_y are the sample means of $\mathbf{x}_T$ and $\mathbf{y}_T$, respectively, to respect the scaling of the data when a single kernel has total unit weight. The $\delta^{(1)}$ starting values repeat a vector $\mathbf{D}^{(0)}$; we use $\mathbf{D}^{(0)} = (1, 2, \dots, 25)^T$, which will be repeated four times to fill the population of size $n = 100$. This range is reasonable for the hydrology application but would be changed in other contexts. Similarly, the $\sigma^{(1)}$ starting values repeat $\mathbf{S}^{(0)} = (1, 5, 10, 20)^T$. When $k > 1$ and window k is added, $\beta^{(1)}, \dots, \beta^{(k)}$ are all reduced in magnitude to $m_y/m_x/k$ to account for k kernels. The starting values for $\delta^{(1)}$ and $\sigma^{(1)}$ are as for the $k = 1$ iteration; the $\delta^{(k)}$ and $\sigma^{(k)}$ for the further windows are given by a random permutation of the elements in $\mathbf{D}^{(0)}$ and $\mathbf{S}^{(0)}$, respectively. This procedure ensures the starting values for a multi-window iteration have sufficient coverage of the parameter space, while having distinct initial values for each of the k windows.

If the number of windows is not known a priori, an upper bound $k_{\max}$ must be defined, and the hyperparameter $\hat{k}$ will be selected by choosing the best among all $k_{\max}$ iterations based on BIC.

Note that the suggested model initializations are specific to the hydrological application tackled in this work. A different choice of hyperparameters may be required if the model was transferred to another domain. The Gaussian SWR model is implemented in R version 4.3.0 (R Core Team, 2022). The implementation is publicly available on GitHub.¹

2.5 Autocorrelated residuals

We expect autocorrelation to be present in the residuals of the Gaussian SWR model, since both input and output time series may have other effects or measurement errors correlated with time. This aspect violates the requirements of least squares likelihood and may distract evaluation metrics such as mean squared error. Thus, we aim to resolve such issues by adapting the estimation procedure proposed by Cochrane and Orcutt (1949).

Suppose the errors ε_t in the model (4) follow an autoregressive process of lag 1, i.e. AR(1), so that

$$\varepsilon_t = \varphi \varepsilon_{t-1} + \eta_t, \quad (6)$$

where η_t is Gaussian white noise (instead of ε_t in equation (4)), and $\varphi \in (-1, 1)$. From equation (4) and (6), it follows that

$$\begin{aligned} y_t &= \sum_{i=1}^k \beta^{(i)} \cdot (\mathbf{x}_{[t]} * \boldsymbol{\kappa}^{(i)}) + \varphi \varepsilon_{t-1} + \eta_t, \text{ and} \\ y_{t-1} &= \sum_{i=1}^k \beta^{(i)} \cdot (\mathbf{x}_{[t-1]} * \boldsymbol{\kappa}^{(i)}) + \varepsilon_{t-1}. \end{aligned}$$

In order to convert this setup into a model with Gaussian white noise errors, we investigate the expression

$$\begin{aligned} y_t - \varphi y_{t-1} &= \sum_{i=1}^k \beta^{(i)} \cdot (\mathbf{x}_{[t]} * \boldsymbol{\kappa}^{(i)} - \varphi \mathbf{x}_{[t-1]} * \boldsymbol{\kappa}^{(i)}) + \eta_t \\ &= \sum_{i=1}^k \beta^{(i)} \cdot \left(\sum_{s=1}^t (x_s - \varphi x_{s-1}) \boldsymbol{\kappa}_{t-s}^{(i)} \right) + \eta_t. \end{aligned}$$

¹ <https://www.github.com/sschranner/SlidingWindowReg>.

The term $x_0 \cdot \kappa_t^{(i)}$ originating from index $s = 0$ is ignored since $\kappa^{(i)}$ is padded with 0 at the end (when being adjusted to the length of $\mathbf{x}_{[t]}$). Hence, given φ , the model can be written as an ordinary least squares model with uncorrelated errors η_t :

$$\hat{y}_t = \sum_{i=1}^k \beta^{(i)} \cdot (\tilde{\mathbf{x}}_{[t]} * \kappa^{(i)}) + \eta_t,$$

where the transformation $\tilde{z}_t = z_t - \varphi z_{t-1}$ can be applied a priori as part of the preprocessing procedure on both variables x_t and y_t .

For the Gaussian SWR model, we adapt the Cochrane–Orcutt procedure to estimate φ and the optimal model parameters:

- estimate model parameters β , δ , σ from the original model using $\mathbf{x}_{[t]}$ and $\mathbf{y}_{[t]}$ (i.e. set $\varphi = 0$),
- estimate φ from the model residuals,
- apply the transformation $\tilde{z}_t = z_t - \varphi z_{t-1}$ to both $\mathbf{x}_T$ and $\mathbf{y}_T$ and obtain final estimates of the model parameters β , δ , σ .

The same concept also holds for more general AR(m)-processes with $m > 1$. In general, the transformation leads to a loss of m time points (samples) in the training dataset.

2.6 Uncertainty quantification

Minimizing the negative log-likelihood with respect to (β, δ, σ) gives maximum-likelihood estimates of the model parameters. The large-sample theory implies that the errors in the estimators have an approximate multivariate Gaussian distribution with mean 0 (asymptotically unbiased) and an approximate covariance matrix $\mathcal{I}^{-1}$, where

$$\mathcal{I}(\beta, \delta, \sigma) = \frac{\partial^2}{\partial(\beta, \delta, \sigma)^2} (-\ln L(\beta, \delta, \sigma; \mathbf{x}_T, \mathbf{y}_T))$$

denotes the observed information, i.e. the negative second derivative of the log-likelihood evaluated at the MLE parameter values. The Hessian can be obtained numerically via the R package `numDeriv` (Gilbert & Varadhan, 2019). The diagonal elements of $\mathcal{I}^{-1}$ lead to approximate standard errors, which provide the uncertainty quantification reported below for the two watersheds studied. These watersheds have a very long time series to support the asymptotic argument.

3 Simulation study

In a simulation study, the proposed Gaussian SWR model will be validated in controlled scenarios based on real-world input data and simulated targets.

3.1 Experimental setup

Precipitation data $\mathbf{x}_T$ is used from the Koksilah River watershed which is located in Cowichan, British Columbia, Canada, while an artificial target variable $\mathbf{y}_T$ is sampled based on equation (4). Both time series are acquired on a daily basis. The model parameters β , δ , and σ are randomly generated for this purpose, along with Gaussian white noise errors $\varepsilon_t \sim N(0, \rho^2)$. The distributions of the model parameters and level of the error variance ρ^2 are described further below.

In all of the following setups, the model parameters of the Gaussian SWR model were estimated on a training set comprising a time series of the first 29 hydrological years, resulting in 10,593 data points. The test set consisted of the remaining 10 time series of hydrological years or 3,652 data points, which corresponds to an approximate 75%/25% split of the dataset.

The primary factors steering the difficulty associated with a setup in the simulation study are:

1. the (ground truth) number of windows k^{gt} —the dimensionality and complexity of parameter estimation increases with the number of windows,

2. the pairwise overlap between (ground truth) windows on the time axis—a higher overlap leads to a decrease in the separability of windows, and
3. the level of measurement noise ρ^2 applied, given by the noise rate (relative to the explained variance of the model, as highlighted below).

While k^{gt} and ρ are systematically varied, the overlap of windows is a result from the random sampling of ground truth model parameters.

We employ 15 distinct setups (five 1-, 2-, and 3-window setups, respectively). A higher number of windows is not common for hydrological models due to the limited number of possible flow paths. Each setup uses a set of parameters independently sampled from $\delta^{(i)} \sim U(0, 20)$, and $\sigma^{(i)}, \beta^i \sim U(1, 5)$ for any $i \in \{1, \dots, k\}$, where $U(a, b)$ denotes a uniform distribution on the interval $[a, b]$. Each 1-, 2-, and 3-window setup is repeated at five distinct levels of measurement noise $\alpha \in \{0.05, 0.25, 0.5, 0.75, 0.95\}$ to control the error relative to the signal, for a total of 75 simulation experiments.

Given a noise level α , the measurement noise is added as follows. After evaluating the deterministic terms in equation (4) with the given model parameters, denoted as $\tilde{\mathbf{y}}_T$, we compute the associated sample variance $\text{Var}(\tilde{\mathbf{y}}_T)$. Gaussian white noise with standard deviation $\rho = \alpha \cdot \sqrt{\text{Var}(\tilde{\mathbf{y}}_T)}$ is simulated and added to construct $\mathbf{y}_T = \tilde{\mathbf{y}}_T + \boldsymbol{\varepsilon}_T$.

The resulting sample variance of $\mathbf{y}_T$ can be expressed as $\text{Var}(\mathbf{y}_T) = (1 + \alpha^2)\text{Var}(\tilde{\mathbf{y}}_T)$. Thus, α^2 represents the relative noise level with respect to the explained variance. This induces an upper bound on the R^2 score (Kuhn & Johnson, 2019) achievable by any predictor given $\mathbf{x}_T$ and $\mathbf{y}_T$: a perfect fit giving $\hat{\mathbf{y}}_t = \tilde{\mathbf{y}}_T$ would achieve an R^2 score of

$$\begin{aligned}
 R^2(\mathbf{y}_T, \hat{\mathbf{y}}_T) &= 1 - \frac{\|\mathbf{y}_T - \hat{\mathbf{y}}_T\|_2^2}{\|\mathbf{y}_T - \bar{\mathbf{y}}\mathbf{1}\|_2^2} \\
 &= 1 - \frac{\text{Var}(\boldsymbol{\varepsilon}_T)}{\text{Var}(\mathbf{y}_T)} \\
 &= 1 - \frac{\alpha^2 \text{Var}(\tilde{\mathbf{y}}_T)}{(1 + \alpha^2) \text{Var}(\tilde{\mathbf{y}}_T)} \\
 &= \frac{1}{1 + \alpha^2},
 \end{aligned} \tag{7}$$

where $\bar{\mathbf{y}}$ denotes the mean of $\mathbf{y}_T$ across T , and $\mathbf{1}$ denotes a vector of ones. All noise levels α and associated R^2 scores are shown in Table 1.

3.2 Evaluation metrics

In order to evaluate the Gaussian SWR model, we deploy evaluation metrics to assess both the accuracy of parameter estimation and the predictive performance. To facilitate the assessment of the

Table 1. Noise levels and associated maximum R^2 scores (upper bound for predictive model performance)

	Noise level α	Maximum R^2
1	0.05	0.998
2	0.25	0.941
3	0.5	0.8
4	0.75	0.64
5	0.95	0.526

estimated parameters, we introduce the notion of *overlap* between two weight vectors with equal dimensions, $\mathbf{w}^{(1)}, \mathbf{w}^{(2)} \in \mathbb{R}^t$ as follows:

$$O(\mathbf{w}^{(1)}, \mathbf{w}^{(2)}) = \sum_{s=1}^t \min \{w_s^{(1)}, w_s^{(2)}\}.$$

Hence, in a functional setting, the overlap would describe the area covered by the minimum of both kernel functions. Since both weight vectors are positive and sum to 1 due to the normalization condition, the overlap will be in $[0, 1]$ and reaches 1 only if both weight vectors are exactly equal. For a multiple-window setting, a combined kernel takes account of the kernel weights but is then normalized for ease of comparison when computing the overlap.

We group the evaluation criteria applied to our model results into two categories:

- *Kernel overlap*: When combining all estimated windows to a joint weight vector

$$\sum_{i=1}^k \beta^i \kappa^{(i)},$$

i.e. the mixture of all kernels, the overlap between the full predicted and the ground truth kernels is assessed (including the weighting of windows by regression parameters β). In particular, we compute the overlap $O(\hat{\theta}, \theta)$, where $\hat{\theta} = (\hat{\beta}, \hat{\delta}, \hat{\sigma})$, and $\theta = (\beta, \delta, \sigma)$, respectively. Thus, the combined kernel overlap evaluates the quality of overall parameter estimation, including the regression parameters β .

- *Predictive performance*: Conventional regression metrics are used to evaluate the predictive performance of the model on the test set (Kuhn & Johnson, 2019). These include the root mean squared error (RMSE), as well as the coefficient of determination (R^2) on the test set. The latter is also known as Nash–Sutcliffe Efficiency (NSE) in hydrology (Nash & Sutcliffe, 1970) and is introduced in equation (7). In addition, the Kling–Gupta efficiency (KGE) is a performance metric used in hydrology which, unlike the NSE, independently encourages predictions to match the variability of observations, thereby removing the tendency of NSE to underestimate high flows and overestimate low flows (Gupta et al., 2009). It uses the following formulation:

$$1 - \sqrt{(r-1)^2 + \left(\frac{\sigma_{\hat{y}_T}}{\sigma_{y_T}} - 1\right)^2 + \left(\frac{\mu_{\hat{y}_T}}{\mu_{y_T}} - 1\right)^2},$$

where r denotes the Pearson correlation coefficient between $\mathbf{y}_T$ and $\hat{\mathbf{y}}_T$, μ represents the respective means, and σ represents the respective standard deviations. All performance metrics are implemented using the R package `hydroGOF` (Zambrano-Bigiarini, 2020).

3.3 Results for uncorrelated errors

First, we visualize the predicted kernels obtained from training the Gaussian SWR model in each setup specified above. In the experiment, the algorithm selects the number of windows k based on the best BIC, subject to a maximum of $k_{\max} = 3$ windows. This upper limit originates from our application, where typically a maximum number of up to three distinct flow paths (overland, sub-surface, groundwater flow) is assumed.

Figure 6 compares the weights at discrete time points of the ground truth and estimated kernels for the middle noise level $\alpha = 0.5$ in each setup. Dotted vertical lines indicate the true or estimated window centres $\delta^{(i)}$ along the time axis. In all 1-, 2-, and 3-window setups, window positions and sizes were accurately predicted. Moreover, dominant peaks and general shapes are reconstructed in all scenarios. Since we allow noninteger numbers for delta, the discretized weights are not always symmetric w.r.t. the measured time points. Overall, the combined kernel matches with the ground truth very accurately, and the general separation of the windows is acceptable.




Figure 6. Weights at discrete time points of the ground truth versus predicted kernels (noise level $\alpha = 0.5$). The number of windows was selected as a model hyperparameter.

Beyond visual comparison, Table 2 provides a summary of window overlap accuracy based on predicted and true parameters. Overall, all setups achieve excellent overlaps with the ground truth kernels. Thus, the reconstruction of the overall weighted kernel is at a high level across all simulated setups.

Having observed the high accuracy of parameter identification in the simulation experiment, we further investigate the predictive performance of the models with respect to the target variable on the test set. Figure 7 illustrates the corresponding R^2 , KGE, and RMSE scores obtained by the estimated Gaussian SWR models. In agreement with the observations made on the quality of parameter estimation, the predictive performance remains at a high level and is mainly affected by the noise levels. For the R^2 score, the upper bounds in Table 1 are indicated by a shaded region, which cannot be reached at the given noise level. The achieved average R^2 values are close to the upper bounds across all model setups, and hence the estimated models obtain almost optimal prediction accuracy on the test set.

To guarantee interpretability of the model parameters, accurate estimation of the true number of windows k^{gt} is crucial. For k^{gt} equal to 1, 2, or 3, all 75 estimated models select the hyperparameter k correctly.

Table 2. Parameter estimation accuracy for 1-, 2-, and 3-window setups across five noise levels

Setup no.	Noise level α				
	0.05	0.25	0.5	0.75	0.95
1-Window setups					
1	1.00	1.00	1.00	1.00	0.99
2	1.00	1.00	1.00	1.00	0.99
3	1.00	1.00	1.00	1.00	1.00
4	1.00	1.00	1.00	0.99	1.00
5	1.00	1.00	0.99	1.00	0.99
Mean	1.00	1.00	1.00	1.00	0.99
2-Window setups					
6	1.00	1.00	0.99	0.99	0.99
7	1.00	1.00	0.99	0.99	0.99
8	1.00	1.00	0.99	0.99	0.99
9	1.00	1.00	0.99	0.99	0.98
10	1.00	1.00	1.00	0.99	0.98
Mean	1.00	1.00	0.99	0.99	0.99
3-Window setups					
11	1.00	0.99	0.99	0.98	0.98
12	0.99	0.99	0.98	0.98	0.97
13	1.00	1.00	0.99	1.00	0.97
14	1.00	0.98	0.99	0.98	0.98
15	1.00	0.99	0.98	0.97	0.96
Mean	1.00	0.99	0.99	0.98	0.97

Note. An overlap of 1 indicates perfect agreement between an estimated window and the ground truth.

3.4 Results for autocorrelated errors

To account for a more general error structure, we repeat the above experiment but with response data simulated with autocorrelated AR(1) errors ε_t . We make use of the R function `arma.sim` from the `forecast` package (Hyndman & Athanasopoulos, 2021) to simulate AR(1) noise with a given standard deviation ρ . Analogous to the case of uncorrelated errors, ρ is based on the sample variance $\text{Var}(\tilde{y}_T)$ and the noise levels α from Table 1. Internally, the function constructs i.i.d. errors η_t with the given standard deviation ρ , which are then transformed to form an AR(1) process. The true autoregressive parameter φ is set to 0.5 throughout for the purpose of demonstrating the concept.

In practice, the autocorrelation has to be assessed, and the autoregressive parameter φ needs to be estimated when we extend the Gaussian SWR model. The Durbin–Watson test for autocorrelated model errors is performed to assess the residuals before and after the Cochrane–Orcutt data transformation. As expected, model fits with no Cochrane–Orcutt transformation deliver p -values below .01 (rejecting the null hypothesis of no autocorrelation) for all 75 setups. The average absolute estimation error of the parameter φ is less than 0.006 across all setups, suggesting highly accurate estimation. After applying the Cochrane–Orcutt procedure, 74 out of the 75 setups lead to residuals with p -values above .05, which provides no evidence for violation of the assumption of uncorrelated residuals in the model fit to transformed data. The remaining setup, with p -value of .032, has only small residual autocorrelation of 0.02.

With autocorrelated errors, the upper bounds for the R^2 scores achievable by the model must be modified. From (6), we have $\text{Var}(\varepsilon_t) = \text{Var}(\varphi\varepsilon_{t-1} + \eta_t) = \varphi^2 \text{Var}(\varepsilon_{t-1}) + \text{Var}(\eta_t)$ as η_t is uncorrelated




Figure 7. Predictive performance summaries for 1-, 2-, and 3-window setups by noise level α . All results are averaged over the five distinct simulated parameter setups.

with ε_{t-1} by construction. Iterating this result, and noting that $\text{Var}(\varepsilon_1) = \text{Var}(\eta_1)$ and $\text{Var}(\eta_t) = \alpha^2 \text{Var}(\tilde{y}_T)$, gives

$$\begin{aligned}
 \text{Var}(\varepsilon_t) &= \varphi^2 \text{Var}(\varepsilon_{t-1}) + \text{Var}(\eta_t) \\
 &= \varphi^4 \text{Var}(\varepsilon_{t-2}) + \varphi^2 \text{Var}(\eta_{t-1}) + \text{Var}(\eta_t) \\
 &= \sum_{i=0}^{t-1} \varphi^{2i} \text{Var}(\eta_{t-i}) \\
 &= \underbrace{\frac{1 - \varphi^{2(t-1)}}{1 - \varphi^2}}_{\zeta} \alpha^2 \text{Var}(\tilde{y}_T).
 \end{aligned}$$

It holds that $\zeta \approx \frac{1}{1-\varphi^2}$ for longer time series. As a result, the maximum achievable R^2 score (when $\hat{y}_T = \tilde{y}_T$) under the AR(1) model errors is given by

$$\begin{aligned}
 R^2(y_T, \hat{y}_T) &= 1 - \frac{\zeta \alpha^2 \text{Var}(\tilde{y}_T)}{(1 + \zeta \alpha^2) \text{Var}(\tilde{y}_T)} \\
 &= \frac{1}{1 + \zeta \alpha^2}.
 \end{aligned}$$

Similar to the simulation with uncorrelated residuals, the estimated kernels closely align with the ground truth, demonstrating accurate reconstruction of the overall weighted kernel across all simulated setups. For the given noise levels α , the maximum R^2 scores for the autocorrelated simulations are presented in Table 3. Results obtained for the experimental setups are shown in Figure 8. As with the uncorrelated model setups, the R^2 scores almost reach the theoretical upper bound.

Table 3. Noise levels and associated maximum R^2 scores (upper bound for predictive model performance) in the autocorrelated setups

	Noise level α	Maximum R^2
1	0.05	0.997
2	0.25	0.923
3	0.5	0.75
4	0.75	0.571
5	0.95	0.454




Figure 8. Predictive performance summaries for 1-, 2-, and 3-window setups by noise level α in the autocorrelated setup.

Overall, the hyperparameter k representing the number of selected windows, determined via BIC matches the ground truth in all 1-window setups, overestimates by 1 for 3 out of 25 of the 2-window setups, and underestimates by 1 for one of the 25 3-window simulations.

4 Real-world data experiments

In the second part of our experiments, we evaluate the performance of the Gaussian SWR model on two real-world examples. Both watersheds are located on the west coast of North America, with the first catchment representing the Koksilah River which is located in Cowichan, British Columbia, Canada. The second catchment represents Big Sur River in central California, United States. The available dataset consists of 39 hydrological years for both watersheds. Both datasets are available on the aforementioned GitHub repository (entitled `sampleWatershed` and `sampleWatershed2`, respectively).

4.1 Experimental setup

Analogous to the simulation study, we use a training subset containing years 1–29 of the time series, followed by a test set covering 10 years. The maximum number of windows is set to $k_{\max} = 4$.

After initial training, the model residuals are evaluated for autocorrelation using the Durbin-Watson test. If autocorrelation is detected, the Cochrane–Orcutt procedure is applied to correct for AR(1) errors. If the Durbin–Watson test finds no evidence of autocorrelation in the updated model the process stops; otherwise, AR(2) errors are considered, etc.

4.2 Results

The Durbin–Watson test on the residuals from the model with untransformed data gives a p -value $< .01$ in both cases, i.e. strong evidence against an independent-error model. In the case of watershed 1, an AR(2) process with parameters $\phi_1 = 0.46$ and $\phi_2 = 0.12$ is indicated to model the autocorrelation, resulting in a Durbin–Watson p -value of 0.20 after applying the Cochrane–Orcutt procedure. For watershed 2, an AR(1) process with parameter $\phi = 0.84$ appears to be sufficient to remove the autocorrelation, resulting in a p -value of .39 after data transformation.

The estimated kernels for each watershed are illustrated in Figure 9, and their parameters are shown along with uncertainties in Table 4. Predictive performance scores (R^2 , KGE , and $RMSE$) on the test set are shown in Table 5. Overall, the Gaussian SWR model achieves reasonably accurate predictions for both watersheds, e.g. R^2 scores of approximately 0.71 and 0.6, respectively, on the untransformed streamflow scale.

The Gaussian SWR model selects $k = 3$ windows for both Koksilah and Big Sur Rivers. Figure 9a shows a clear dominance of the short-term history (a window close to a lag of 1 day) for the Koksilah River. Paired with the fact that only three highly overlapping windows were needed to model the catchment with high accuracy (Table 5), we can conclude that the Koksilah River is driven by simple hydrological processes. The third window is very wide with lower weight, and the first two windows have $\delta \simeq 1$, thus the combined kernel for watershed 1 resembles a truncated Student- t distribution: the combined kernel remains unimodal, but the probability mass is more dispersed over the long-term lags compared to a Gaussian kernel. Moreover, a low lag value suggests a relatively simple catchment, which aligns with the observation that the Koksilah River is situated in a region characterized by extremely high precipitation and minimal aridity (Janssen & Ameli, 2021). From previous works, we can hypothesize that wet, nonarid, catchments tend to be extremely simple with fast response times (Schoppa et al., 2020; Spieler et al., 2020), which is strengthened by our experimental results. Since the bulk of precipitation transitions into streamflow in about 1 day, this catchment is likely dominated by shallow subsurface flow. Large amounts of precipitation in this catchment fully saturate and connect the shallow soil layers for much of the year, resulting in any excess precipitation pushing the water already in the soils towards the stream within one day.

As we travel south along the western coast of North America towards the Big Sur River in Central California, aridity increases and total precipitation decreases slightly (Janssen & Ameli, 2021). Although the Big Sur River is still considered a ‘wet’ catchment, the Gaussian SWR model selects two more distinguishable peaks for watershed 2 compared to watershed 1, indicating an increase in catchment complexity in-part due to the increased aridity (Schoppa et al., 2020;




Figure 9. Estimated weights at discrete time points of scaled window kernels in real-world data experiments. The combined kernel (dashed line) for watershed 1 resembles a truncated Student-*t* distribution with stronger tails than Gaussian probability density functions. The combined kernel (dashed line) for watershed 2 is multimodal. (a) Koksilah River (Watershed 1). (b) Big Sur River (Watershed 2).

Table 4. Parameter estimates after Cochrane–Orcutt transformation for real-world watersheds, with approximate standard errors

Watershed no	Window no	$\beta^{(i)}$	$\delta^{(i)}$	$\sigma^{(i)}$
1	1	0.39 ± 0.01	1.24 ± 0.22	0.18 ± 0.12
	2	0.35 ± 0.02	1.28 ± 0.46	2.34 ± 0.44
	3	0.17 ± 0.02	7.31 ± 9.09	18.12 ± 0.65
2	1	0.21 ± 0.01	0.00 ± 0.33	0.56 ± 0.24
	2	0.31 ± 0.02	2.00 ± 0.01	2.67 ± 0.02
	3	0.40 ± 0.02	14.79 ± 0.65	10.40 ± 0.42

[Spieler et al., 2020](#)). The first large peak centred at $\delta = 0$ indicates overland flow. This peak could be the result of the Big Sur River flowing through several developed towns where any precipitation falling on impermeable urban surfaces causes immediate drainage into the river. The wider second

Table 5. Predictive performance metrics (on test set) after parameter estimation using Cochrane–Orcutt transformation for real-world watersheds

Watershed	k	Metric		
		R^2	KGE	RMSE
Original data				
1	3	0.71	0.72	3.36
2	3	0.60	0.70	2.74
Transformed data				
1	2	0.59	0.64	2.67
2	3	0.47	0.58	1.50

Note. Evaluation metrics are computed on the original (untransformed) test set, as well as on the transformed test set. KGE, Kling-Gupta efficiency; RMSE, root mean squared error.




Figure 10. Observed streamflow, predicted streamflow, and rainfall in the real-world experiments for six years of the test set. (a) Watershed 1. (b) Watershed 2.

Table 6. Runtimes of the Gaussian SWR model with $k_{\max} \in \{1, 2, 3\}$, averaged over five runs

$k_{\max}$	Runtime [s]
1	12.4
2	75.6
3	608.8

window could represent shallow subsurface flow in a similar fashion as watershed 1. The last window for the Big Sur River indicates that there is substantial baseflow in the catchment. With some precipitation flowing deeper underground towards less porous material, this partition of water will respond more slowly to water input compared to shallow subsurface flow and overland flow.

Figure 10 compares the time series of predicted and observed water flows for each watershed, along with the respective rainfall time series (all relating to untransformed data). To save space, the plots restrict the time domain to the first three and last three hydrological years of the test set. Since the Gaussian SWR model acts as a smoothing operator on the input time series, the prediction time series is less spiked than the rainfall input and generally matches the shape of the ground truth well. Particularly for watershed 1, time intervals with flat gauge values appear more noisy in the predictions. This becomes visible in the dry periods at the end of each hydrological year (time interval [250, 365]), where the ground truth response is almost exactly 0 throughout, while the predictions are modestly positive.

To investigate the effect of sample size, we repeated training for Watershed 1 but with only the first 6 years of data, i.e. about 20% of the previous training data comprising 29 years. The previously found third, very wide window with smaller weight seen in Figure 9a was no longer detected by BIC using the smaller dataset, but the combined kernel from the two windows closely resembled that shown in Figure 9a from three windows.

4.3 Runtime analysis

All real-world experiments were performed on an Intel Core i5 3 GHz CPU with 8GB RAM running macOS Sonoma 14.4, where Algorithm 1 was implemented in R. The runtimes in Table 6 are averages over five independent model fits to account for variations in the computation time. In each run, the Gaussian SWR model is trained for one real-world watershed using the same experimental setup as in the real-world study; time to compute predictions is excluded. The maximum number of windows $k_{\max}$ is clearly important, but all runtimes are moderate.

5 Discussion and conclusion

Our experiments validated the Gaussian SWR model in a controlled experimental environment and real-world scenarios. Good predictive accuracy on a test set, as measured by R^2 , for instance, was achieved throughout. The simulations demonstrated that achievable accuracy was dominated by the noise level, and rather independent of the number, position, size, and weighting of the ground truth windows. Moreover, multiple kernels were usually identifiable, with accurate parameter estimates. Thus, the proposed Gaussian SWR model has promising utility for hydrological inference.

Parameterizing the kernels as having Gaussian densities is easily extended to more flexible shapes, such as Student's t , Gamma, or asymmetric generalized Laplace distributions. The latter two suggestions would enable long-tailed lag effects.

Relative to the more general DLM, the predictive performance of the Gaussian SWR model is comparable on the datasets used in our real-world experiments. When varying the number of time lags q , the DLM implementation by Demirhan (2020) achieves test-set R^2 scores of 0.7 ($q = 5$), 0.72 ($q = 10$ and $q = 20$) for watershed 1 and 0.51 ($q = 5$), 0.55 ($q = 10$), 0.58 ($q = 20$) for watershed 2. Thus, the parametric assumption taken by Gaussian SWR, i.e. restricting the kernel shape to a mixture of Gaussians, does not significantly affect the predictive performance. However, the favourable parametrization employed by the Gaussian SWR model allows for straightforward interpretations of the model parameters: the window position and shape parameters are directly transferable to the concept of distinct flow paths in hydrology and deliver

information on the peak lag of the runoff and its celerity distribution, respectively. Regression parameters β give indications of the relative importance of each flow path: while a dominant window with short lag typically models a strong short-term (overland) flow, well-distinguished windows with longer time lags indicate the presence of distinct long-term effects such as subsurface or groundwater flows. From a statistical perspective, kernels representing long-term effects may approximate a constant line, thus adding an intercept could be a valid option for the proposed model—for the hydrological interpretation, however, a large (slow-flow) window acts as a more accurate parameterization than a model with intercept.

From a hydrological viewpoint, a major limitation of the Gaussian SWR model is that the kernels, particularly the parameters β and δ , do not change dynamically over the hydrologic year. The hydrologic year in many watersheds is characterized by drier and wetter periods, with distinct patterns in rainfall–runoff relationships. Our concurrent work on the same watershed 1 demonstrates that the predictive performance can be increased to approximately $R^2 = 0.8$ using a functional data analysis approach with dynamic regression weights and sparsity. Potentially, the Gaussian SWR model setup could be extended in future work to account for such dynamic parameters at the cost of a higher number of model parameters and a more complex training procedure, while keeping its highly interpretable ability to separate flow paths.

The presence of significant precipitation falling as snow will impact runoff dynamics in colder regions. As a consequence, runoff lags and magnitudes of the distinct flows paths may vary significantly over the year, which is not covered by linear constant-lag models. To overcome this issue for watersheds with significant snow-to-rainfall ratios, an extension of the proposed Gaussian SWR model to account for nonlinearities would further be of interest. We consider this as another potential future extension of the proposed model.

Acknowledgments

We would like to thank Asad Haris for contributing to the preparation of the dataset, as well as Oliver Meng and Junsong Tang for their work on graphical user interfaces to visualize the proposed model.

Conflicts of interest: No conflict of interest is declared.

Funding

This research was funded by the Canadian Statistical Sciences Institute (CANSSI) as a Collaborative Research Team project. S.S. and A.J.I gratefully acknowledge financial support from the Norwegian University of Life Sciences (project number 1211130114) to visit the University of British Columbia. J.C.'s research is supported by the Canada Research Chairs program and the Discovery grant (RGPIN-2023-04057) of the Natural Sciences and Engineering Research Council of Canada (NSERC). Computing was supported in part by the Digital Research Alliance of Canada (alliancecan.ca).

Data availability

Our code and data are available from the following GitHub repository: <https://www.github.com/sschranner/SlidingWindowReg> (entitled `sampleWatershed` and `sampleWatershed2`, respectively).

Author contributions

S.S., A.J., and W.J.W. proposed the method, J.J. prepared the dataset, S.S., A.J., and P.P. implemented the method and conducted the experiment(s), J.J., J.C., and P.P. suggested improvements to the method, J.J. and A.A.A. provided hydrological background, S.S. and J.J. wrote the first draft of the manuscript, and all authors reviewed and revised the manuscript.

References

- Ahmad S. K., & Hossain F. (2020). Maximizing energy production from hydropower dams using short-term weather forecasts. *Renewable Energy*, 146, 1560–1577. <https://doi.org/10.1016/j.renene.2019.07.126>

- Almon S. (1965). The distributed lag between capital appropriations and expenditures. *Econometrica: Journal of the Econometric Society*, 33(1), 178–196. <https://doi.org/10.2307/1911894>
- Baltagi B. H. (2022). *Econometrics* (6th ed). Springer Nature.
- Chen Y.-H., Mukherjee B., & Berrocal V. J. (2018). Distributed lag interaction models with two pollutants. *Journal of the Royal Statistical Society: Series C, Applied Statistics*, 68(1), 79–97. <https://doi.org/10.1111/rssc.12297>
- Cochrane D., & Orcutt G. H. (1949). Application of least squares regression to relationships containing autocorrelated error terms. *Journal of the American Statistical Association*, 44(245), 32–61. <https://doi.org/10.2307/2280349>
- Cornette N., Roques C., Boisson A., Courtois Q., Marçais J., Launay J., Pajot G., Habets F., & de Dreuz J.-R. (2022). Hillslope-scale exploration of the relative contribution of base flow, seepage flow and overland flow to streamflow dynamics. *Journal of Hydrology*, 610, Article 127992. <https://doi.org/10.1016/j.jhydrol.2022.127992>
- Davtyan A., Rodin A., Muchnik I., & Romashkin A. (2020). Oil production forecast models based on sliding window regression. *Journal of Petroleum Science & Engineering*, 195, 107916. <https://doi.org/10.1016/j.petrol.2020.107916>
- Demirhan H. (2020). dLagM: An R package for distributed lag models and ARDL bounds testing. *PLoS One*, 15(2), e0228812. <https://doi.org/10.1371/journal.pone.0228812>
- Dunn S., Birkel C., Tetzlaff D., & Soulsby C. (2010). Transit time distributions of a conceptual model: Their characteristics and sensitivities. *Hydrological Processes*, 24(12), 1719–1729. <https://doi.org/10.1002/hyp.v24:12>
- Eisner R. (1960). A distributed lag investment function. *Econometrica: Journal of the Econometric Society*, 28(1), 1–29. <https://doi.org/10.2307/1905291>
- Everitt B. S., Hand D. J., & Everitt B. (1981). *Finite mixture distributions*. Monographs on Statistics & Applied Probability. Chapman and Hall.
- Giani G., Rico-Ramirez M. A., & Woods R. A. (2021). A practical, objective, and robust technique to directly estimate catchment response time. *Water Resources Research*, 57(2), e2020WR028201. <https://doi.org/10.1029/2020WR028201>
- Gilbert P., & Varadhan R. (2019). numDeriv: Accurate numerical derivatives. R package version 2016.8-1.1. <https://CRAN.R-project.org/package=numDeriv>
- Griliches Z. (1967). Distributed lags: A survey. *Econometrica: Journal of the Econometric Society*, 35(1), 16–49. <https://doi.org/10.2307/1909382>
- Gupta H. V., Kling H., Yilmaz K. K., & Martinez G. F. (2009). Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling. *Journal of Hydrology*, 377(1-2), 80–91. <https://doi.org/10.1016/j.jhydrol.2009.08.003>
- Hannan E. J. (1967). The estimation of a lagged regression relation. *Biometrika*, 54(3-4), 409–418. <https://doi.org/10.1093/biomet/54.3-4.409>
- Hare D. K., Helton A. M., Johnson Z. C., Lane J. W., & Briggs M. A. (2021). Continental-scale analysis of shallow and deep groundwater contributions to streams. *Nature Communications*, 12(1), Article number 1450. <https://doi.org/10.1038/s41467-021-21651-0>
- Harman C., Troch P., & Sivapalan M. (2011). Functional model of water balance variability at the catchment scale: 2. Elasticity of fast and slow runoff components to precipitation change in the continental United States. *Water Resources Research*, 47(2), 1–12. <https://doi.org/10.1029/2010WR009656>
- Hyndman R. J., & Athanasopoulos G. (2021). *Forecasting* (3rd ed). OTexts.
- Janssen J., & Ameli A. A. (2021). A hydrologic functional approach for improving large-sample hydrology performance in poorly gauged regions. *Water Resources Research*, 57(9), e2021WR030263. <https://doi.org/10.1029/2021WR030263>
- Janssen J., Radić V., & Ameli A. (2021). Assessment of future risks of seasonal municipal water shortages across North America. *Frontiers in Earth Science*, 9, 730631. <https://doi.org/10.3389/feart.2021.730631>
- Kannan N., White S., Worrall F., & Whelan M. (2007). Hydrological modelling of a small catchment using SWAT-2000—Ensuring correct flow partitioning for contaminant modelling. *Journal of Hydrology*, 334(1-2), 64–72. <https://doi.org/10.1016/j.jhydrol.2006.09.030>
- Khan I. A., Akber A., & Xu Y. (2019). Sliding window regression based short-term load forecasting of a multi-area power system. In *2019 IEEE Canadian Conference of Electrical and Computer Engineering (CCECE)*. IEEE. <https://doi.org/10.1109/ccece.2019.8861915>
- Kuhn M., & Johnson K. (2019). *Applied predictive modeling*. Springer.
- McMillan H. (2020). Linking hydrologic signatures to hydrologic processes: A review. *Hydrological Processes*, 34(6), 1393–1409. <https://doi.org/10.1002/hyp.v34.6>
- Mebane W. R., & Sekhon J. S. (2011). Genetic optimization using derivatives: The rgenoud package for R. *Journal of Statistical Software*, 42(11), 1–26. <https://doi.org/10.18637/jss.v042.i11>
- Mebane W. R., Sekhon J. S., & Saarinen T. (2023). rgenoud: R version of GENetic optimization using derivatives. <https://10.32614/CRAN.package.rgenoud>

- Nash J., & Sutcliffe J. (1970). River flow forecasting through conceptual models. Part I. A discussion of principles. *Journal of Hydrology*, 10(3), 282–290. [https://doi.org/10.1016/0022-1694\(70\)90255-6](https://doi.org/10.1016/0022-1694(70)90255-6)
- Nejadhashemi A. P., Shirmohammadi A., Sheridan J. M., Montas H. J., & Mankin K. R. (2009). Case study: Evaluation of streamflow partitioning methods. *Journal of Irrigation and Drainage Engineering*, 135(6), 791–801. [https://doi.org/10.1061/\(ASCE\)IR.1943-4774.0000093](https://doi.org/10.1061/(ASCE)IR.1943-4774.0000093)
- Peng R. D., Dominici F., & Welty L. J. (2009). A Bayesian hierarchical distributed lag model for estimating the time course of risk of hospitalization associated with particulate matter air pollution. *Journal of the Royal Statistical Society: Series C, Applied Statistics*, 58(1), 3–24. <https://doi.org/10.1111/j.1467-9876.2008.00640.x>
- Powell M. J. (2009). *The BOBYQA algorithm for bound constrained optimization without derivatives* (Cambridge NA Rep. NA2009/06). University of Cambridge, 26.
- R Core Team (2022). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- Roksvåg T., Steinsland I., & Engeland K. (2021). A two-field geostatistical model combining point and areal observations—A case study of annual runoff predictions in the Voss area. *Journal of the Royal Statistical Society: Series C, Applied Statistics*, 70(4), 934–960. <https://doi.org/10.1111/rssc.12492>
- Rushworth A. M., Bowman A. W., Brewer M. J., & Langan S. J. (2013). Distributed lag models for hydrological data. *Biometrics*, 69(2), 537–544. <https://doi.org/10.1111/biom.v69.2>
- Sawicz K., Kelleher C., Wagener T., Troch P., Sivapalan M., & Carrillo G. (2014). Characterizing hydrologic change through catchment classification. *Hydrology and Earth System Sciences*, 18(1), 273–285. <https://doi.org/10.5194/hess-18-273-2014>
- Schoppa L., Disse M., & Bachmair S. (2020). Evaluating the performance of random forest for large-scale flood discharge simulation. *Journal of Hydrology*, 590, 125531. <https://doi.org/10.1016/j.jhydrol.2020.125531>
- Schwarz G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6(2), 461–464. <https://doi.org/10.1214/aos/1176344136>
- Spieler D., Mai J., Craig J. R., Tolson B. A., & Schütze N. (2020). Automatic model structure identification for conceptual hydrologic models. *Water Resources Research*, 56(9), e2019WR027009. <https://doi.org/10.1029/2019WR027009>
- Tang Q., Rosenberg E., & Lettenmaier D. (2009). Use of satellite data to assess the impacts of irrigation withdrawals on Upper Klamath Lake, Oregon. *Hydrology and Earth System Sciences*, 13(5), 617–627. <https://doi.org/10.5194/hess-13-617-2009>
- Wang Y.-H., & Gupta H. V. (2024). Towards interpretable physical-conceptual catchment-scale hydrological modeling using the mass-conserving-perceptron. *Water Resource Research*, 60(10), e2024WR037224. <https://doi.org/10.1029/2024WR037224>
- Warren J. L., Luben T. J., & Chang H. H. (2020). A spatially varying distributed lag model with application to an air pollution and term low birth weight study. *Journal of the Royal Statistical Society: Series C, Applied Statistics*, 69(3), 681–696. <https://doi.org/10.1111/rssc.12407>
- Young P. C. (2006). *Rainfall-Runoff modeling: Transfer function models*. John Wiley & Sons, Ltd. Chapter 128. <https://doi.org/10.1002/0470848944.hsa141a>
- Zalewski M. (2000). Ecohydrology—The scientific background to use ecosystem properties as management tools toward sustainability of water resources. *Ecological Engineering*, 16(1), 1–8. [https://doi.org/10.1016/S0925-8574\(00\)00071-9](https://doi.org/10.1016/S0925-8574(00)00071-9)
- Zambrano-Bigiarini M. (2020). *hydroGOF: Goodness-of-fit functions for comparison of simulated and observed hydrological time series*. <https://doi.org/10.32614/CRAN.package.hydroGOF>